Causal state-dependent local projections

- Identification and inference in R
- An applied workshop for DG ECFIN
- Weining Wang | University of Bristol

Two policy applications

- Countries: institutional reforms and economic performance
- Firms: financial conditions and innovative activity
- Report patterns, uncertainty and the scope of causal claims

A flexible three-hour workshop

- 100 minutes: essential concepts and guided R work
- 20 minutes: breathing room for questions and timing drift
- 20 minutes: breaks; 10 minutes: short discussions
- Final 30 minutes: open questions and research clinic

Define the causal comparison first

- Specify outcome, horizon, shock units and initial state
- Change the date-t innovation; let later outcomes evolve
- Compare responses across states without changing the states

$$\mathrm{IRF}_h(z;\delta) = \delta g_h(z), \qquad \Delta_h = g_h(z_A) - g_h(z_B)$$

Get the timing right

- State and controls: information available before the shock
- Shock: date t; outcome: date t+h
- Respect publication dates, anticipation and data revisions

Shock construction is an identification decision

- A policy change need not be an exogenous policy shock
- Assess anticipation, information effects and omitted news
- Residualizing a series does not establish causal exogeneity
- An instrument requires an IV design and IV inference

The identification conditions

- Causal exogeneity links shocks to potential outcomes
- An unpredictable shock removes the past-information nuisance
- The conditional mean is affine in the shock
- The state captures the relevant response coefficient; support is positive

$$E[X_t \mid \mathcal{F}_{t-1}] = 0, \qquad E[Y_{i,t+h} \mid \mathcal{F}_{t-1}, X_t] = g_h(Z_{i,t-1})X_t + r_{i,h}(\mathcal{F}_{t-1})$$

Why the moment identifies the response

- Multiply the conditional mean by the shock
- The predictable nuisance vanishes under the MDS condition
- Divide by conditional shock variation

$$g_h(z) = \frac{E[X_t Y_{i,t+h} \mid Z_{i,t-1} = z]}{E[X_t^2 \mid Z_{i,t-1} = z]}$$

An endogenous state can still be admissible

- Initial conditions may depend on earlier shocks
- Stable propagation preserves linearity in the current shock
- Changing propagation or shock nonlinearities need more care

$$Y_{it} = a(Z_{i,t-1})X_t + \rho Y_{i,t-1} + \varepsilon_{it} \quad \Longrightarrow \quad g_h(z) = \rho^h a(z)$$

State choice can change the weighting

- A coarse state groups units with different true responses
- LP weights can depend on conditional shock variance
- An ordinary subgroup average needs additional conditions

$$\frac{1+3}{2} = 2, \qquad \frac{1 \cdot 1 + 4 \cdot 3}{1+4} = 2.6$$

Discuss your identification design

- Name your outcome, shock, initial state and horizon
- Identify one plausible threat to causal interpretation
- Decide whether the current aim is descriptive or causal

Questions on identification

- Clarify the causal target and timing
- Keep unresolved application questions for the final clinic

Break

- Ten minutes
- After the break: open the R scripts and simulated panel

The simulated panel

- 100 units observed over 240 shock dates
- A common shock; a pre-shock financial-state index
- Known response: a hump in the state, with stable decay
- Synthetic units: no empirical policy calibration

$$g_h(z) = 0.6^h(1 - z^2)$$

Start in R

- Extract the package and set the working directory
- 01_course.R: transparent matrix implementation
- 02_paper_estimator.R: B-splines, AIC and inference

```
source("R/01_course.R")
source("R/02_paper_estimator.R")
```

The linear interaction is a benchmark

- Estimate the common shock and its interaction with the state
- A small interaction does not establish homogeneous effects
- Compare the fitted function with a flexible alternative

$$Y_{i,t+h} = \alpha_h X_t + \beta_h X_t Z_{i,t-1} + W_{i,t-1}^{\top} \gamma_h + u_{i,t+h}$$

A flexible response with cubic B-splines

- Interact every basis column with the shock
- Keep the same basis when predicting at new state values
- Choose J from 4 to 20 using the paper's AIC rule

$$\widehat{g}_h(z) = \phi_J(z)^\top \widehat{b}_h, \qquad D_{it} = (X_t \phi_J(Z_{i,t-1})^\top, W_{i,t-1}^\top)^\top$$

```
B <- splines::bs(z, degree=3, knots=knots,
                 intercept=TRUE)
D <- cbind(B*x, W)
```

R lab: check timing and the sample

- Check unique unit-date keys and missing periods
- Verify that z is pre-shock and x is common within date
- Inspect state support and the number of shock dates

```
stopifnot(!anyDuplicated(d[c("id","t")]))
quantile(d$z, c(.05,.5,.95))
length(unique(d$t))
```

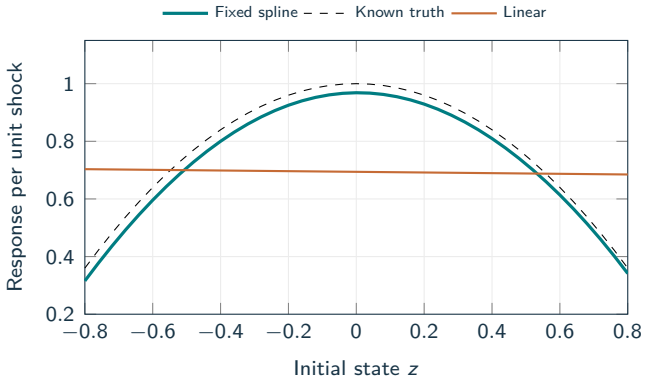
R lab: estimate the response function

- Run the fixed-basis comparison, then the paper estimator
- Inspect the AIC table and selected basis dimension
- Plot the response only over the chosen data-supported range

```
impact <- estimate_horizon(0)
impact$AIC
impact$J
infer_contrast(impact)
```

The linear fit misses the hump

- Synthetic impact response
- The flexible fit recovers substantial state variation
- The linear interaction is nearly flat



Inference starts with the score process

- Sum regression scores across units within each date
- Estimate serial dependence in that date-level process
- Sandwich the long-run covariance with the OLS bread
- Use state-specific standard errors

$$s_t = \sum_i D_{it} \widehat{u}_{i,t+h}, \qquad \widehat{V} = (D^\top D)^{-1} \widehat{\Omega} (D^\top D)^{-1}$$

```
scores <- rowsum(D*u, time)
bread <- solve(crossprod(D))
V <- bread %*% meat %*% bread
```

R lab: inference for a state contrast

- Choose the states before looking for a large difference
- Use the covariance between the two fitted responses
- Report the estimate, interval, horizon and shock units

$$a = \phi(z_A) - \phi(z_B), \qquad \widehat{\Delta} = a^\top \widehat{b}, \qquad \mathrm{se}(\widehat{\Delta}) = \sqrt{a^\top \widehat{V}_b a}$$

```
infer_contrast(impact, zA=0, zB=.75)
```

Discuss the R output

- Which result would you put in an internal note?
- Which additional evidence would a journal paper need?

Questions and coding buffer

- Time for questions arising during the practical
- Optional: fixed J, alternative HAC lags, or common samples

Break

- Ten minutes
- Next: simultaneous bands, pitfalls and your applications

Responses across horizons

- Rebuild the outcome and estimation sample at each horizon
- Distinguish levels, changes and cumulative outcomes
- Pointwise horizon intervals do not cover the whole path jointly

```
fits <- lapply(0:8, estimate_horizon)
```

Pointwise intervals and simultaneous bands

- Pointwise: one preselected state at one horizon
- Simultaneous band: the whole specified state grid at one horizon
- Studentize Gaussian draws by each state's standard error
- A state band is not a band over every horizon

$$c_{.95} = q_{.95}\left[\max_{z \in \mathcal{Z}_G} \frac{|\phi(z)^\top \widehat{V}_b^{1/2} \xi|}{\widehat{\sigma}(z)}\right], \qquad \widehat{g}(z) \pm c_{.95}\widehat{\sigma}(z)$$

```
curve <- infer_curve(impact,
  grid=seq(-.8,.8,length.out=161))
```

Diagnostics before interpretation

- Count independent shock dates, not only panel rows
- Inspect support, shock variation and influential dates
- Check missingness, sample composition and long horizons
- Inspect sensitivity to state definition and spline complexity

Controls and fixed effects need a purpose

- Use pre-shock information for the intended total response
- Post-shock controls can remove transmission channels
- Date fixed effects absorb the common shock main effect
- Unit fixed effects bring their own dynamic-panel issues

A valid shock does not guarantee the target

- Nonlinear shock effects can break the affine restriction
- A shock can change persistence or the propagation mechanism
- A diagnostic rejection is informative; non-rejection is not proof

$$Y = X^2, \quad X \text{ symmetric} : \quad \frac{E[XY]}{E[X^2]} = 0, \quad Y(\delta) - Y(0) = \delta^2$$

Adapt the R code to your data

- Map the outcome, shock, state, controls and calendar explicitly
- Reassess identification for this shock and data structure
- Change the dependence model when the design changes
- Keep transformations, samples and tuning choices reproducible

Write a defensible result

- State the target, shock units, horizon and population
- Name the uncertainty: pointwise, contrast, or simultaneous
- Separate empirical patterns from causal and policy claims
- Explain the limitation that could change the conclusion

Replicate first, then assess sensitivity

- Match the published target, sample, transformations and units
- Match controls, horizons, tuning and covariance conventions
- Reproduce a reported number before changing the design
- Then test sensitivity to choices that matter for the claim

```
estimate_horizon(4, Jset=7)
estimate_horizon(4, bandwidth=8)
```

Before using an LP result for policy

- Can I defend the shock and initial-state timing?
- Does the response-function interpretation fit this setting?
- Does inference match dependence and the scope of the claim?
- Have I reported support, units and meaningful sensitivity?

Final clarification buffer

- Resolve remaining questions about the core workflow
- Move extended research discussions into the clinic

Open questions and research clinic

- Describe your data, outcome, shock and initial state
- What would make the result useful for your project?
- Identify the main identification or inference difficulty
- Agree one concrete next analysis or diagnostic

Research clinic: next steps

- Each participant names one next step
- Identify any question needing a different estimator or design
- Keep the R scripts as a starting point for documented adaptation

Optional case: institutional reforms

- Country panel: reform timing, anticipation and reverse causality
- Persistent indicators, bundled reforms and cross-border spillovers
- Country and time effects; dependence across countries and years
- A pooled response need not apply to each member state

Optional case: firm finance and innovation

- Define financing shock separately from financial condition
- Date the initial state and innovation outcome carefully
- Consider patent lags, zeros, selection and firm exit
- Match dependence to the shock's level of variation

Optional: Causal identification: the first half of the proof

- Consistency and causal exogeneity identify mean potential outcomes
- The affine mean implies the same slope at each admissible shock value
- Subtracting two potential-outcome means removes the nuisance

$$E[Y_{i,t+h}(x) \mid \mathcal{F}_{t-1}] = g_h(Z_{i,t-1})x + r_{i,h}(\mathcal{F}_{t-1}),$$
$$\mathrm{IRF}_h(z;\delta) = \delta g_h(z).$$

Optional: Statistical identification: the shock moment

- Multiply the conditional mean by the shock
- The MDS condition removes the additive nuisance
- Condition on the predetermined state
- Divide by the conditional shock second moment

$$E[X_t Y_{i,t+h} \mid Z = z] = g_h(z) E[X_t^2 \mid Z = z],$$
$$g_h(z) = \frac{E[X_t Y_{i,t+h} \mid Z = z]}{E[X_t^2 \mid Z = z]}.$$

Optional: The slope averages local derivatives

- For a smooth response, the interaction coefficient averages derivatives
- The derivative weights are nonnegative and integrate to one
- A positive average derivative does not establish global monotonicity
- A small coefficient can reflect cancellation of opposite slopes

$$\beta_h = \int_{\mathcal{Z}} \omega_h(z) g_h'(z)\, dz, \qquad \omega_h(z) \geq 0, \qquad \int_{\mathcal{Z}} \omega_h(z)\, dz = 1$$

Optional: Shock construction, residualisation and instruments

- A residualised shock must remove the relevant predictable component
- Residualisation alone does not establish causal exogeneity
- A proxy instrument needs relevance, exclusion and a clear normalization
- Generated shocks and weak instruments require additional inference

$$\widetilde{X}_t = X_t - E[X_t \mid \mathcal{F}_{t-1}]$$

Optional: Common shocks and fixed effects

- The aggregate shock is identical across units at a date
- Unrestricted date effects absorb its common level
- Shock interactions may retain cross-sectional variation
- Absolute response levels then require a normalization or additional structure

$$X_t \in \text{span}\{\text{date indicators}\}$$

Optional: Identification and uncertainty in micro-macro panels

- A large cross-section does not create more aggregate shock dates
- Common LP errors can survive cross-sectional averaging
- Overlapping horizons produce serial dependence
- Pointwise intervals and uniform bands answer different questions

Optional: Appendix: weighted and unweighted state targets

- The ordinary subpopulation response averages causal slopes
- The LP additionally weights by conditional shock variance
- Their difference is a conditional covariance divided by mean variance
- Constant conditional variance is sufficient for equality

$$\frac{\mathrm{IRF}_h(z;\delta)}{\delta} = E[\widetilde{g}_h(s) \mid Z = z],$$

$$g_h^{\mathrm{LP}}(z) - E[\widetilde{g}_h(s) \mid Z = z] = \frac{\mathrm{Cov}(v(s,z), \widetilde{g}_h(s) \mid Z = z)}{E[v(s,z) \mid Z = z]}, \quad v(s,z) = E[X^2 \mid s, Z = z].$$